

PCT

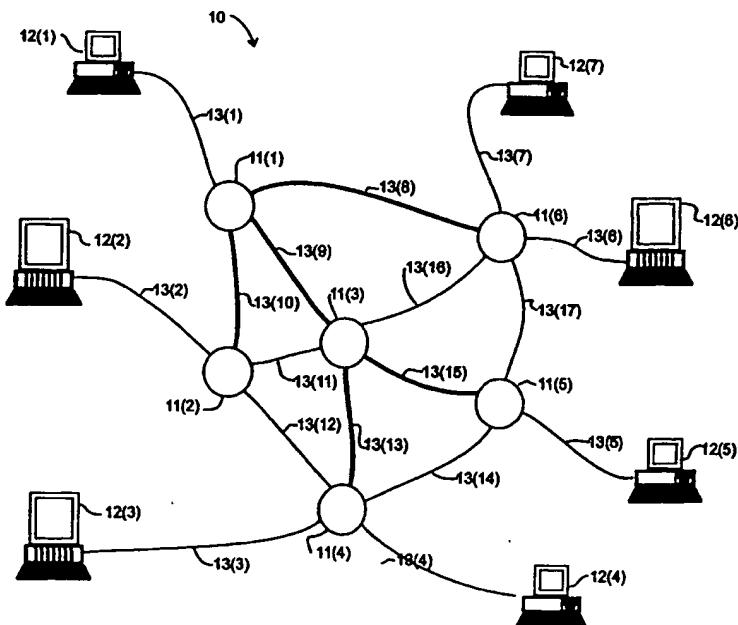
WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ :	A1	(11) International Publication Number:	WO 98/18239
H04L 12/46, 12/56		(43) International Publication Date:	30 April 1998 (30.04.98)
(21) International Application Number:	PCT/US97/19558		
(22) International Filing Date:	24 October 1997 (24.10.97)		
(30) Priority Data:	<p>08/738,411 25 October 1996 (25.10.96) US</p> <p>(71) Applicant: CASCADE COMMUNICATIONS CORPORATION [US/US]; 5 Carlisle Road, Westford, MA 01886 (US).</p> <p>(72) Inventor: MOY, John, T.; 8 Castle Drive, Acton, MA 01720 (US).</p> <p>(74) Agent: JORDAN, Richard, A.; P.O. Box 81363, Wellesley Hills, MA 02181-0004 (US).</p>		
(81) Designated States:	AU, CA, JP, European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).		
<p>Published <i>With international search report. Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i></p>			

(54) Title: MULTIPONT-TO-POINT PACKET TRANSFER OVER VIRTUAL CIRCUIT



(57) Abstract

A network (10) comprises a plurality of switching nodes (11(1)-11(6)) interconnected by a plurality of communication links (13(8)-13(16)). At least one virtual circuit is established in the network (10) to at least one terminating switching node (11(n)) defining a path which the other switching nodes (11(n)) can use to transfer message packets to the terminating switching node (11(n)).

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republie of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

MULTIPOINT-TO-POINT PACKET TRANSFER OVER VIRTUAL CIRCUIT**FIELD OF THE INVENTION**

The invention relates generally to the field of digital networking systems, and more particularly to systems and methods for controlling the transfer of message packets over a digital network using a virtual circuit packet transfer methodology. The invention provides an arrangement in which each virtual circuit can comprise one or more transfer paths through the network all of which terminate at a destination device in a network.

BACKGROUND OF THE INVENTION

Digital networks have been developed to facilitate the transfer of information, including data and programs, among digital computer systems and numerous other types of devices. A variety of types of networks have been developed and implemented using diverse information transfer methodologies. In modern networks, information is transferred through a mesh of switching nodes which are interconnected by communication links in a variety of patterns. The mesh interconnection pattern can allow for a number of paths to be available through the network from each computer system or other device which may transmit information as a source device, to another computer system or other device, which is to receive the information as a destination device, so that if congestion develops in particular regions of the network, or if components of the network become inoperative, information can be routed around the congested or inoperative portion of the network.

Information transferred from a source device to a destination device is generally transferred in the form of fixed- or variable-length packets, which are transferred through the network in accordance with one of several general packet transfer methodologies. In one packet transfer methodology, termed a "connectionless" data transfer methodology, each information packet that is transmitted by a source device to a destination device includes, in addition to the information to be transferred, destination address information which identifies the particular device that is to receive the packet, along with an identifier identifying the particular switching node to which the destination device is connected. When a switching node receives a packet that is to be transferred to a particular destination device, it (that is, the switching node) will use a routing table to identify a communication link over which the packet is to be transmitted and transmit the packet over the identified communication link. Each switching node that receives the packet will perform a similar operation, until the packet is transferred to a switching node which is connected to the destination device. At that point, the switching node will identify the particular communication link connected

-2-

between it and the destination device, and will transfer the packet to the destination device. Similar operations occur for each packet transferred through the network, so that each switching node will make a routing decision for each packet that it receives to identify a particular communication link over which it is to transfer the packet. The routing decision is overhead that may delay delivery of packets to their respective destinations.

Another packet transfer methodology reduces this overhead. In this packet transfer methodology, packets are transferred through the network's switching nodes over constructs which are generally termed "virtual circuits," "virtual connections" and the like (generally, "virtual circuits"). When a source device wishes to transfer information to a destination device, initially a preliminary operation will be performed to establish a virtual circuit over communication links defining a path from the source device, through a series of switching nodes to the destination device. Each switching node in the virtual circuit methodology is provided with a virtual circuit database that associates each of the virtual circuits maintained by the respective switching node with a particular one of the communication links.

After a virtual circuit is established, the source device can transfer packets using the virtual circuit database. Each packet includes, as well as the data to be transferred, an identifier that identifies the particular virtual circuit to be used in transferring the packet. The source device will initially transfer the packet to the switching node connected thereto, and that switching node will use virtual circuit identifier and the virtual circuit database to identify a communication link over which it (that is, the switching node) is to transfer the packet. These operations will be performed for each switching node along the path of the virtual circuit. When the last switching node along the virtual circuit receives the packet, that switching node will transfer the packet to the particular destination device associated with the virtual circuit. These operations will be repeated for each of the packets transferred from the source device to the destination device. In addition, similar operations may also be performed in connection with packets which the destination device may transfer to the source device, since a single virtual circuit, once established, may be used to transfer packets in both directions between the original source device and the original destination device.

At the end of a communications session between the source device and the destination device, during which one or a number of packets may be transferred, the devices may close the virtual circuit. In that operation, packets are transferred between the source and destination devices which, among other things, can enable the switching nodes along the virtual circuit to expunge the virtual circuit information relating to the particular virtual circuit being closed from their respective virtual

-3-

circuit databases. Thereafter, the virtual circuit identifier and other resources, including the entries in the respective virtual circuit databases associated with the closed virtual circuit, can be used for other virtual circuits.

There are several benefits of using each of the aforementioned packet transfer methodologies which are outlined in the aforementioned Chen application and will not be repeated herein. The aforementioned Chen application describes a network which provides, in a network constructed defined according to the "connectionless" packet transfer methodology, a number of the advantages that are provided by the virtual circuit methodology. In particular, the network described in the Chen application comprises a plurality of switching nodes which are interconnected by communication links, and the switching nodes themselves, rather than devices connected thereto, establish virtual circuits over which "connectionless" message packets generated the devices will be transferred. The virtual circuits in the network described in the Chen application are not established and removed at the beginning and end of a communications session, but instead are relatively permanent, and so at least some of the overhead typically associated with a network constructed in connection with the virtual circuit transfer methodology will be avoided. However, each of the switching nodes maintains virtual circuit information relating to all of the virtual circuits for which it is the originating, terminating or an intermediate node along the path defined by the virtual circuit. Since the amount of virtual circuit information that is maintained by the switching nodes will increase on the order of " N^2 " where " N " corresponds to the number of switching nodes in the network, in a network comprising a large number of switching nodes the amount of virtual circuit information required to be maintained can be quite large.

SUMMARY OF THE INVENTION

The invention provides a new and improved system and method for controlling the transfer of message packets over a digital network using a virtual circuit packet transfer methodology in which each virtual circuit can comprise one or more transfer paths through the network.

In brief summary, the invention in one aspect provides a network comprising a plurality of switching nodes interconnected by a plurality of communication links. The switching nodes transfer message packets from an information source to an information destination. At least one virtual circuit is established in said network to at least one "terminating" switching node, the virtual circuit extending over a plurality of other switching nodes in the network. The virtual circuit defines a path which the other switching nodes can use to transfer message packets to the "terminating" switching node.

In another aspect, the invention provides a network comprising a plurality of switching nodes interconnected by a plurality of communication links. The switching nodes transfer message packets from an information source to an information destination. At least one virtual circuit is established in said network to at least one "terminating" switching node, the virtual circuit defining a plurality of paths extending to a like plurality of other switching nodes in the network. The other switching nodes can use the respective paths to transfer message packets to the "terminating" switching node.

The invention effectively allows a number of nodes to use the same virtual circuit for transferring message packets to the "terminating" node, effectively providing a "multi-point to point" virtual circuit. Since a plurality of nodes can use the same virtual circuit, the amount of virtual circuit information that each switching node needs to maintain can be reduced.

BRIEF DESCRIPTION OF THE DRAWINGS

This invention is pointed out with particularity in the appended claims. The above and further advantages of this invention may be better understood by referring to the following description taken in conjunction with the accompanying drawings, in which:

FIG. 1 schematically depicts a computer network constructed in accordance with the invention;

FIGS. 2 and 3 depict structures of several types of message packets used in the network depicted in FIG. 1, which is useful in understanding the invention;

FIG. 4 depicts a functional block diagram of a switching node useful in the network depicted in FIG. 1; and

FIG. 5 is a flow chart depicting operations performed by the switching node depicted in FIG. 4, in connection with the invention.

DETAILED DESCRIPTION OF AN ILLUSTRATIVE EMBODIMENT

FIG. 1 schematically depicts a computer network 10 including a plurality of switching nodes 11(1) through 11(N) (generally identified by reference numeral 11(n)) for transferring signals representing data among a number of devices, which in FIG. 1 are represented by computer systems 12(1) through 12(M) (generally identified by reference numeral 12(m)). The computer systems 12(m), as is conventional, process data, in accordance with their program instructions to generate processed data. In their processing, a computer system 12(m_s) (subscript "S" referencing "source") may, as a source computer system, need to transfer data, processed data and/or program instructions (all of which will be referred to herein generally as "information") to another, destination, computer

system 12(m_D) (subscript "D" referencing "destination"), which may use the transferred information in its operations.

Each computer system 12(m) is connected over a communication link, generally identified by reference numeral 13(p), to a switching node 11(n) to facilitate transmission of data thereto or the reception of data therefrom. The switching nodes 11(n) are interconnected by communication links, also generally identified by reference numeral 13(p) to facilitate the transfer of information among the respective switching nodes 11(n). The communication links 13(p) may utilize any convenient information transmission medium, including, for example, wires for carrying electrical signals, optical fiber links for carrying optical signals, and so forth. Each communication link 13(p) is preferably bi-directional, allowing the switching nodes 11(n) to transmit and receive signals among each other and with computer systems 12(m) connected thereto over the same link; depending on the particular type of medium selected for respective communication links 13(p), multiple media may be provided for transferring signals in opposite directions.

In one embodiment, the computer systems 12(m), as source computer systems 12(m_S), provide information for transfer, organized in packets which are structured according to the well-known "IP" ("Internet Protocol") transfer methodology. In addition, computer systems 12(m) which, as destination computer systems 12(m_D), are to receive the packets, expect to receive such IP packets. In the IP transfer methodology, packets are normally transferred according to the "connectionless" transfer methodology, in which each packet provided by a source computer system 12(m_S), for transfer to a destination computer system 12(m_D), includes a network address that includes a high-order portion, termed herein a "prefix" that identifies the particular switching node 11(n) to which the destination computer system 12(m_D) is connected, as well as a "suffix" that identifies the particular destination computer system 12(m_D) itself. In the network 10, the various switching nodes 11(n) establish virtual circuits thereamong, over which they transfer the packets received from the computer systems 12(m). In the network 10, the various switching nodes 11(n) establish the virtual circuits thereamong automatically, independently of the receipt of message packets from any of the computer systems 12(m) connected thereto. Each virtual circuit can be used to transfer packets generated by a source computer system 12(m_S) to the terminating switching node 11(n_T), which, in turn, will transfer such packets to the destination computer system. Since the virtual circuits are established in the network automatically and independently of the receipt of message packets from any of the computer systems 12(m) connected thereto, the transfers of packets received from the

computer systems 12(m) will not be delayed by overhead operations required to establish the virtual circuits.

In accordance with invention, to reduce the amount of virtual circuit information which the switching nodes need to maintain to define the virtual circuits which they establish, instead of the switching nodes 11(n) establishing a unique virtual circuit between each respective originating switching node 11(n_O) and terminating switching node 11(n_T), as in the system described in the aforementioned Chen patent application, the switching nodes 11(n) in the network 10 (FIG. 1) establish a single virtual circuit for each terminating switching node 11(n_T), which virtual circuit all of the other switching nodes 11(n') (n'≠n_T) which operate as originating switching nodes will use in transferring message packets to the respective terminating switching node 11(n_T). One such virtual circuit will be established for each of the switching nodes which may operate as a terminating switching node 11(n_T). Accordingly, unlike the virtual circuits described in the aforementioned Chen application, in which each virtual circuit defines a single "point to point" path between one respective originating switching node 11(n_O) and one respective terminating switching node 11(n_T), in the network 10 (FIG. 1) each virtual circuit may define a plurality of paths through the network, all of which extend from a remote switching node 11(n_R), as the most remote switching node along a respective path in the virtual circuit, and terminate at the respective terminating switching node 11(n_T). Each virtual circuit in the network 10 is effectively in the form of a "multi-point to point" tree comprising one or more branches, each branch defining a message packet transfer path through the network 10, with the tree having a root at the terminating switching node 11(n_T) and one or more leaves comprising respective remote switching node(s) 11(n_R).

In addition, in the network 10, all of the switching nodes 11(n) along each branch or path of a virtual circuit, in addition to operating as intermediate switching nodes 11(n_I) to forward packets that they receive for the virtual circuit along a respective path comprising the virtual circuit, can also operate as originating switching nodes 11(n_O). As originating switching nodes 11(n_O), the switching nodes 11(n) can transmit message packets which they receive from computer systems 12(m) which are connected thereto, to the terminating switching node 11(n_T), using the same virtual circuit. This is in further contrast to the virtual circuits in the network described in the aforementioned Chen application. In the network described in the Chen application, it a switching node 11(n) along a particular virtual circuit between an originating switching node 11(n_O) and a terminating switching node 11(n_T) operates as an intermediate switching node 11(n_I) for the virtual circuit, if the switching node 11(n) also can operate as an originating switching node, a separate and distinct virtual circuit

is established between that switching node 11(n) and the terminating switching node 11(n_T) over a portion of the same path as that taken by the particular virtual circuit.

Thus, for example, if computer system 12(1) can operate as a destination computer system, a single virtual circuit will be established which extends to the switching node 11(1) connected thereto, as terminating switching node, from the switching nodes 11(2), 11(4), 11(5) and 11(6), as originating switching nodes, if at least one computer system connected to the respective switching node 11(2), 11(4), 11(5) and 11(6) can operate as source computer system. The virtual circuit established by the terminating switching node 11(1) illustratively covers the paths comprising the communication links 13(p) represented by the thick solid lines in FIG. 1. All of the paths share a common "root" at the terminating switching node 11(1), with all of the other switching nodes which can operate as originating switching node, namely, switching nodes 11(2), and 11(4) through 11(6), comprise leaves for the virtual circuit "tree." The paths comprising the virtual circuit include

(i) one path comprising communication link 13(8) from the switching node 11(6), with switching node 11(6) operating as the remote switching node for the path,

(ii) a second path comprising communication links 13(9) and 13(15) from switching node 11(5) through switching node 11(3), with switching node 11(5) operating as the remote switching node for the path,

(iii) a third path comprising communication links 13(9) and 13(13) from switching node 11(4) also through switching node 11(3), with switching node 11(4) operating as the remote switching node for the path, and

(iv) a fourth path comprising communication links 13(10) from switching node 11(2), with switching node 11(4) operating as the remote switching node for the path.

The virtual circuit so established defines a plurality of paths, as indicated above, all of which terminate at the terminating switching node 11(1), which the other switching nodes 11(2) through 11(6) can use to transfer message packets that they receive from the respective computer systems 12(m) connected thereto, to the switching node 11(1) for delivery to the computer system 12(1) as destination computer system. As shown in FIG. 1, the second and third paths share the communication link 13(9) between switching node 11(3) and 11(1), and so the virtual circuit branches at switching node 11(3) to form separate paths to switching nodes 11(4) and 11(5).

Thus, if, for example, a computer system 12(6) or 12(7) generates a message packet for transfer to the computer system 12(1), it (that is, computer system 12(6) or 12(7)) will initially transfer the message packet to the remote switching node 11(6) for the first path described above,

which switching node also operates as originating switching node for the transfer. The remote switching node 11(6), in turn, will use the previously-established virtual circuit for which the switching node 11(1) is the terminating switching node, to transfer the message packet to the switching node 11(1). After the terminating switching node 11(1) receives the message packet, it (that is, the switching node 11(1)) can provide the message packet to the destination computer system 12(1).

Similarly, if computer system 12(5) generates a message packet for transfer to the computer system 12(1), it (that is, computer system 12(5)) will initially transfer the message packet to the remote switching node 11(5) for the second path described above, which switching node also operates as originating switching node for the transfer. The switching node 11(5), in turn, will use the previously-established virtual circuit for which the switching node is the terminating switching node, to initiate transfer the message packet to the switching node 11(1). In the actual transfer operation, the switching node 11(5) will initially transfer the message packet to the switching node 11(3). The switching node 11(3) will operate as intermediate switching node to transfer the message packet to switching node 11(1). After switching node 11(1) receives the message packet, it (that is, the switching node 11(1)) can provide the message packet to the destination computer system 12(1).

On the other hand, if, for example, a computer system 12(3) or 12(4) generates a message packet for transfer to the computer system 12(1), it (that is, computer system 12(3) or 12(4)) will initially transfer the message packet to the remote switching node 11(4) for the third path described above, which switching node also operates as originating switching node for the transfer. The remote switching node 11(4), in turn, will use the previously-established virtual circuit for which the switching node 11(1) is the terminating switching node, to transfer the message packet to the switching node 11(1). In the actual transfer operation, the switching node 11(4) will initially transfer the message packet to the switching node 11(3). The switching node 11(3) will operate as intermediate switching node to transfer the message packet to switching node 11(1). After switching node 11(1) receives the message packet, it (that is, the switching node 11(1)) can provide the message packet to the destination computer system 12(1).

Finally, if computer system 12(2) generates a message packet for transfer to the computer system 12(1), it (that is, computer system 12(2)) will initially transfer the message packet to the switching node 11(2). The switching node 11(2), in turn, will operate as originating switching node for the transfer, and will use the previously-established virtual circuit for which the switching node 11(1) is the terminating switching node, to transfer the message packet to the switching node 11(1).

-9-

After the terminating switching node 11(1) receives the message packet, it (that is, the switching node 11(1), can provide the message packet to the destination computer system 12(1).

It will be appreciated that the switching nodes 11(2) through 11(6) will use the same virtual circuit to transfer the message packets received from the respective computer systems in the examples described above to the switching node 11(1), that is, the virtual circuit defined by the terminating switching node 11(1). In addition, it will also be appreciated that the switching nodes 11(4) and 11(5) will both use a path through the switching node 11(3) as an intermediate switching node. If a computer system 12(m) (not shown) which can operate as a source computer system 12(m_s), were also connected to the switching node 11(3), that switching node 11(3) could also use the same virtual circuit for message packets provided thereto by the computer system for transfer to the computer system 12(1).

As will be described below in greater detail, each virtual circuit which is maintained by a switching node 11(n) will be defined by virtual circuit information, including, for example, a virtual circuit identifier which identifies the particular virtual circuit and the identification of a particular communication link 13(p) that is connected to the switching node 11(n) over which a message packet received over the communication link is to be transmitted if the switching node is not the terminating switching node for the virtual circuit, or an indicator if the switching node is the terminating switching node for the virtual circuit. It will be appreciated that, by having all of the switching nodes 11(n) ($n \neq 1$) in the network 10 which may operate as originating switching nodes, use the same virtual circuit to transfer message packets to the switching node 11(1), the amount of virtual circuit information each switching node will need to store can be reduced from that required in the network described in the aforementioned Chen application.

This will be apparent from the following. In the network described in the aforementioned Chen application, separate virtual circuits are provided extending from each of the originating switching nodes to each of the terminating switching nodes. Accordingly, in the network described in the aforementioned Chen application, a separate virtual circuit would be provided between each of the switching nodes 11(2), 11(4), 11(5), and 11(6), as respective originating switching nodes, and the switching node 11(1), as terminating switching node. In that case, the switching node 11(3) would need to maintain virtual circuit information for two virtual circuits, namely,

(i) the virtual circuit extending from switching node 11(4) as originating switching node to switching node 11(1) as terminating switching node, to facilitate transfers of message packets from computer systems 12(3) and 12(4) to computer system 12(1), and

-10-

(ii) the virtual circuit extending from switching node 11(5) as originating switching node to switching node 11(1) as terminating switching node, to facilitate transfers of message packets from computer systems 12(5) to computer system 12(1).

In addition, if the switching node 11(3) itself were an originating switching node, which would occur if a source computer system were connected thereto, it (that is, the switching node 11(3)) would also maintain virtual circuit information for a third virtual circuit, namely, a virtual circuit extending from the switching node 11(3) to the switching node 11(1). It will be appreciated that the switching node 11(3) will maintain such virtual circuit information for each virtual circuit for which it is an originating switching node for which any of the switching nodes 11(n) ($n \neq 3$) is the terminating switching node, and for which it is the intermediate switching node between any originating switching node 11(n) and terminating switching node 11(n') ($n' \neq n$). Accordingly, the number of virtual circuits, for which each switching node 11(3) stores virtual circuit information, will be on the order of N^2 , where "N" corresponds to the number of switching nodes 11(n) in the network. Each of the other switching nodes 11(n) ($n \neq 3$) will also store virtual circuit information for similar numbers of virtual circuits. Thus, for networks with large numbers of switching nodes 11(n), the amount of virtual circuit information to be stored can grow significantly.

On the other hand, in the network 10 described herein, since the same virtual circuit is used by all of the switching nodes 11(n') ($n' \neq n$) may operate as originating switching nodes for transmission of message packets to switching node 11(n) as terminating switching node, each switching node need only maintain virtual circuit information for one virtual circuit for each terminating switching node. Thus, there will be on the order of "N" virtual circuits in the network 10, where "N" corresponds to the number of switching nodes 11(n) in the network 10. By allowing all of the switching nodes, as originating switching nodes, to share the same virtual circuit to transfer message packets to each terminating switching node, the number of virtual circuits for which each switching node 11(n) will need to store virtual circuit information which the switching nodes need to store can be reduced from on the order of " N^2 " to being generally on the order of "N." Accordingly, in a network comprising a large number "N" of switching nodes, the network constructed in accordance with the invention can provide a significant reduction in the amount of virtual circuit information which the switching nodes need to store.

With this background, the structure and operation of the network 10 will be described in detail in connection with FIGS. 2 through 5. As indicated above, the respective originating switching nodes 11(n) receive packets which have been generated by the respective source computer

-11-

systems 12(m) for transfer using the "connectionless" IP packet transfer methodology, which packets will be termed herein "device packets." The originating switching nodes 11(n) will generate, from the device packets, "virtual circuit packets" for transfer the packets to respective terminating switching nodes using the virtual circuits as described above. When the terminating switching nodes receive the virtual circuit packets, they will recover the device packets therefrom and provide them to the respective destination computer system 12(m_D). FIGS. 2 and 3 depict the structures of a device packet 20 and a virtual circuit packet 30, respectively, which are useful in one embodiment of the invention.

With initial reference to FIG. 2, a device packet 20 constructed in accordance with the IP packet transfer methodology includes a header portion 21 and a data portion 22. The data portion 22 contains the actual data to be transmitted from the source computer system 12(m_S) to the destination computer system 12(m_D). The header portion 21 includes one or more fields which contain, among other information, routing control information which the switching nodes 11(n) can typically use in routing the packet 20 from the source computer system 12(m_S) to the destination computer system 12(m_D). In particular the header portion 21 includes a destination address field 23, which contains a value corresponding to the network address for the the destination computer system 12(m_D). Generally in the IP methodology, as indicated above, the destination address includes the high-order "prefix" portion that provides an identifier for the node that is connected to the particular destination computer system 12(m_D), as well as the low-order "suffix"portion that identifies the destination computer system 12(m_D) itself.

In addition, the header portion 21 will generally include a source address field 24 and one or more miscellaneous fields generally identified by reference numeral 25. The source address field 24 contains a value corresponding to the network address for the source computer system 12(m_S). The miscellaneous fields 25 may contain a variety of types of protocol information, such as, for example, packet type information identifying the type of packet, which may assist the switching nodes 11(n) and destination computer system 12(m_D) in, for example, identifying the various fields of the header portion 21 and the data portion 22. If the packet can be of variable length, the miscellaneous fields 25 may also include length information identifying the length of the packet. In addition, the miscellaneous fields 25 may contain information which may be used to specifically identify a particular application at the destination computer system 12(m_D) which is to receive and process the data provided in the data portion 22, and the like. It will be appreciated that the sequence of the various fields 23 through 25 in the header 21 may differ from that depicted in FIG. 2, and that

-12-

various ones of the miscellaneous fields 25 may be situated throughout the header portion 21. The device packet 20 may also include an error detection code field (not shown) which includes an error detection code that is useful in verifying that the packet has been properly transferred.

As described above, the switching nodes 11(n) receive the device packets 20 from the various computer systems 12(m) and generate therefrom virtual circuit packets 30 for transfer over the various virtual circuits established in the network 10. FIG. 3 depicts details of a virtual circuit packet 30 useful in connection with the invention. With reference to FIG. 3, virtual circuit packet 30 includes a virtual circuit packet header 31 followed by a portion that comprises the device packet 20, including the device packet's header portion 21 and data portion 22. The virtual circuit packet header 31 comprises fields which contain information which is used by the switching nodes 11(n) transferring the virtual circuit packet 30 over the virtual circuit. In particular, the virtual circuit packet header 31 includes a virtual circuit identifier field 32, which identifies the particular virtual circuit over which the virtual circuit packet 30 is to be transferred. In addition, the virtual circuit packet header 31 may also include one or more virtual circuit packet miscellaneous information fields, generally identified by reference numeral 33, which contain protocol information which is used by the switching nodes 11(n) in transferring the packet 30. The protocol information contained in the miscellaneous information fields 33 may include such information as an identifier identifying the packet 30 as a virtual circuit packet, length information identifying the length of the packet, and other information which may be useful in facilitating the transfer of the packet 30 over the virtual circuit. It will be appreciated that the ordering of the fields 32 and 33 in the header 31 may differ from that depicted in FIG. 3, and that one or more of the miscellaneous fields 33 may be situated ahead of the virtual circuit identifier field 32. The virtual circuit packet 30 may also include an error detection code field (not shown) which includes an error detection code that is useful in verifying that the packet has been properly transferred.

The switching nodes 11(n) comprising network 10 also make use of several other types of message packets, including advertising message packets and virtual circuit path establishment message packets (not shown) in connection with the invention, which are described below.

When a computer system operating as a source computer system 12(m_S) provides a device packet 20 to the switching node 11(n) connected thereto for transfer to a destination computer system 12(m_D), the destination address field 23 will include the network address required for the destination computer system 12(m_D), including the prefix that identifies the switching node 11(n) connected to the destination computer system 12(m_D), as well as the suffix that identifies the destination computer

-13-

system 12(m_D) itself. The switching node 11(n) which receives the device packet from the source computer system 12(m_S) will, as originating switching node 11(n_0):

(1) use the prefix of the network address to identify the virtual circuit between it (that is, the originating switching node 11(n_0) that which receives the device packet) and the switching node 11(n') which, as terminating switching node 11(n_T) is connected to the destination computer system 12(m_D),

(2) embed the device packet 30 in a virtual circuit packet 20, which contains a virtual circuit identifier in the virtual circuit identifier field 32, and

(3) transfer the virtual circuit packet over the communication link that is associated with the particular virtual circuit identified by the virtual circuit identifier in the virtual circuit packet.

Thereafter, each switching node 11(n) which receives the virtual circuit packet, up to the terminating switching node 11(n_T) along the virtual circuit, as intermediate switching node 11(n_i), will use the virtual circuit identifier in the virtual circuit packet to identify the communication link 13(p) connected thereto over which it (that is, the respective switching node 11(n)) will transfer the virtual circuit packet. When the terminating switching node 11(n_T) receives the virtual circuit packet, it (that is, the switching node 11(n)) will extract the original device packet from the virtual circuit packet and use the suffix of the destination network address contained therein to identify the particular destination computer system 12(m_D) that is to receive the device packet. Thereafter, the switching node 11(n_T) can provide the device packet to the destination computer system 12(m_D).

As a specific example, if, for example, computer system 12(5), as source computer system, is to transfer information to computer system 12(1), as destination computer system, it (that is, computer system 12(5)) will generate a device packet 20 including in data portion 22 the information to be transferred and in destination address field 23 a destination address that identifies computer system 12(1) as the destination computer system and switching node 11(1) as the switching node connected to the destination computer system. In addition, computer system 12(5) will provide the device packet to the switching node 11(5) connected thereto. The switching node 11(5), in turn, instead of using the destination address to make a routing decision to identify the next switching node along a path to the destination computer system 12(1) and transferring the device packet 20 to that switching node as is conventional in the connectionless packet transfer methodology, will identify the one of the previously-established virtual circuits to be used for transferring information to the terminating switching node 11(n_T), namely, switching node 11(1), that is connected to the destination computer system 12(1). In addition, the switching node 11(5) will generate a virtual

-14-

circuit packet 30 for the device packet 20, embed the device packet 20 in the virtual circuit packet 30, and provide in virtual circuit identifier field 32 a virtual circuit identifier identifying the virtual circuit. It will be appreciated that the virtual circuit whose identifier will be provided by switching node 11(5) in the virtual circuit packet 30 in this example, is the virtual circuit for which switching node 11(1) is the terminating switching node. In addition, the originating switching node 11(5) will transmit the virtual circuit packet 30 over the communication link 13(p) connected thereto. In this example, the path of the virtual circuit for terminating switching node 11(1) from switching node 11(5) to switching node 11(1) includes switching node 11(3) as intermediate switching node, and so the switching node 11(5) will transmit the virtual circuit packet over communication link 13(15) to the switching node 11(3).

When the switching node 11(3) receives the virtual circuit packet from switching node 11(5), it will use the virtual circuit identifier in virtual circuit identifier field 32 of the virtual circuit packet to identify one of the communication links 13(p) connected thereto over which it is to transfer the virtual circuit packet to the next switching node along the virtual circuit, in this case, communication link 13(9). When switching node 11(1) receives the virtual circuit packet, since it is the terminating switching node 11(n_T) for the virtual circuit, it will extract the device packet from the virtual circuit packet, use the destination address in the device packet to identify the destination computer system 12(1) to receive the device packet, and transfer the device packet over the communication link 13(1) to that computer system. It will be appreciated that, since the switching nodes 11(5) and 11(3) use the pre-established path of the virtual circuit to the terminating switching node 11(1) in transferring the virtual circuit packet, they do not need to make respective routing path decisions when they receive the virtual circuit packet, which can result in enhanced efficiency in the transfer of information from the source computer system 12(5) to the destination computer system 12(1).

Similar operations will be performed in connection with transfers of packets from each of the respective computer systems 12(2) through 12(7), as source computer systems 12(m_S), to computer system 12(1) as destination computer system 12(m_D). It will be appreciated that, in accordance with the invention, for such transfers, the originating switching nodes 11(2) and 11(4) through 11(6) will generate, from the device packets 20 which they receive from the respective source computer systems 12(2) through 12(7), virtual circuit packets 30 that identify the same virtual circuit, namely, the virtual circuit associated with the switching node 11(1) as the terminating switching node 11(n_T). Thus, virtual circuit packets 30 generated by the switching nodes 11(2) and 11(4) through 11(6), for transfer to the switching node 11(1), will all identify the same virtual circuit. Each of the switching

-15-

nodes 11(n) which can operate as an originating switching node 11(n_o), will perform similar operations in connection with generation of virtual circuit packets for transfer to each of the other switching nodes 11(n') (n'≠n) as terminating switching nodes.

A functional block diagram of a switching node 11(n) useful in connection with the invention is depicted in FIG. 4. With reference to FIG. 4, the switching node 11(n) includes a packet processing and switching fabric subsystem 40 which operates under the control of a control subsystem 41. The packet processing and switching fabric subsystem 40 receives packets from ones of the communication links 13(p) connected thereto as input communication links generally identified by reference numeral 42. The packet processing and switching fabric subsystem 40 may receive device packets 20 (received from computer systems 12(m) which are connected thereto), and virtual circuit packets 30 (received from other switching nodes 11(n') which are connected thereto), as described above in connection with FIGS. 2 and 3. In addition, the packet processing and software fabric subsystem 40 may receive various types of network control packets, including virtual circuit control packets which are used to control establishment and extension of virtual circuits through the network 10. In connection with the invention, under control of the control subsystem 41:

(i) if the packet processing and switching fabric subsystem 40 receives a device packet 20, it will embed the device packet 20 in a virtual circuit packet 30, using virtual circuit identifier and miscellaneous information for fields 32 and 33 the header portion 31 provided by the control subsystem 41, and transmit the virtual circuit packet 30 over one of the output communication links 43 associated with the virtual circuit, as identified by the control subsystem 41, thereby to forward the virtual circuit packet over a path of the virtual circuit toward the terminating switching node for the virtual circuit;

(ii) if the packet processing and switching fabric subsystem 40 receives a virtual circuit packet 30:

(a) if the switching node 11(n) is an intermediate switching node along the virtual circuit, it will transmit the virtual circuit packet 30 over one of the output communication links 43 associated with the virtual circuit, as identified by the control subsystem 41, thereby to forward the virtual circuit packet over a path of the virtual circuit toward the terminating switching node for the virtual circuit; but

(b) if the switching node is the terminating switching node 11(n_T) along the virtual circuit, it will process the virtual circuit packet 30 to extract the device packet 20

-16-

embedded therein and transmit the extracted device packet 20 over one of the output communication links 43, as identified by the control subsystem 41, to the destination computer system 12(m_D);

(iii) if the packet processing and switching fabric subsystem 40 receives an advertising message packet, a virtual circuit path establishment message packet, or another type of virtual circuit control packet, from an input communication link 42, it will pass the information in the packet to the control subsystem 41 for processing; and

(iv) if the packet processing and switching fabric subsystem 40 receives an advertising message packet, a virtual circuit path establishment message packet, or another type of virtual circuit control packet, from the control subsystem 41, it will transfer the packet over an output communication link 43 as indicated in the respective packet.

Details of the structure and operation of packet processing and switching fabric 40 useful in one embodiment of the invention are described in the aforementioned Chen patent application, and will not be described in detail herein.

The control subsystem 41 makes use of several tables for controlling the operations of the packet processing and switching fabric 40, including a routing table 44, a virtual circuit table 45, and a network topology database 46. Details of the structures of the routing table 44 and virtual circuit table 45 are described in the aforementioned Chen application and will not be described in detail herein. In addition, the structure of the network topology database 46 is conventional and will not be described in detail herein.

Generally, the routing table 44 contains information which associates each switching node 11(n'), by the prefix of the switching node 11(n'), with a particular virtual circuit and a particular output communication link 13(p). In addition, the routing table 44 identifies the series of communication links 13(p₁), 13(p₂)...13(p_T) from the switching node 11(n) to the switching node 11(n'). In particular, the routing table 44 associates particular prefix values with particular virtual circuits identified in the virtual circuit table 45. The control subsystem 41 uses this information when it is constructing a virtual circuit packet 30 from a device packet 20, to identify the particular virtual circuit identifier to be used in the virtual circuit packet 30. After the virtual circuit packet 30 is constructed, the control subsystem 41 can use the virtual circuit table 45 to identify the particular one of the output communication links 43 over which the virtual circuit packet is to be transmitted. It will be appreciated that the prefix value identifies the terminating switching node 11(n) connected to the destination computer system 12(mD) that is to receive the device packet 20. Accordingly, in

connection with the invention, for each prefix value, the virtual circuit tables 44 of all of the switching nodes 11(n) in the network 10 will identify the same virtual circuit, although, as will be described below, the particular values of the virtual circuit identifiers which are used to identify the virtual circuit may differ as among the respective switching nodes 11(n).

The virtual circuit table 45 contains information for each of the various virtual circuits in which the switching node 11(n) is included. The control subsystem 41 uses the information in the virtual circuit table 45 to identify a particular one of the output communication links 43 over which the virtual circuit packet is to be transmitted, for each virtual circuit packet 30 generated by the switching node 11(n) as originating switching node 11(n_O) or received by the switching node 11(n) from an input communication link 42 as an intermediate switching node 11(n_I) along the path to the terminating switching node 11(n_T) for the virtual circuit. In addition, the virtual circuit table 45 indicates whether the switching node 11(n) is the terminating switching node for the virtual circuit.

In one particular embodiment, each switching node 11(n) can actually make use of several virtual circuit identifier values which identify each virtual circuit, including:

- (i) a "received" virtual circuit identifier value for virtual circuit packets that are received from another switching node 11(n') (n'≠n), and
- (ii) a "transmitted" virtual circuit identifier value that is used in transmitting the virtual circuit packet over one of the output communication links 43 to another switching node 11(n") (n"≠n', n), and the virtual circuit table 45 will identify each of the virtual circuit identifier values along with the identification of the output communication link 43 over which the virtual circuit packet 30 is to be transmitted.

For each switching node 11(n) which can operate as a terminating switching node 11(n_T):

- (a) If an switching node 11(n') (n'≠n) is an originating switching node, the routing table 44 of the switching node 11(n') will identify the transmitted virtual circuit identifier value (item (ii) above) for the virtual circuit associated with the switching node 11(n), as well as the communication link over which virtual circuit packets associated with that virtual circuit are to be transmitted.
- (b) If the switching node 11(n') is an intermediate switching node for the same virtual circuit, its virtual circuit table 45 will contain at least one entry which associates

- (I) the received virtual circuit identifier value (item (i) above) for virtual circuit packets that it (that is, the switching node 11(n')) receives associated with the virtual circuit, to

-18-

- (II) the transmitted virtual circuit identifier value (item (ii) above) for virtual circuit packets that it (that is, the switching node 11(n') associated with the virtual circuit that it is to transmit.

The virtual circuit table 45 of the switching node 11(n') preferably has one such entry for each communication link 13(p) which forms a part of the path for the virtual circuit over which it can receive a virtual circuit packet for the virtual circuit. Thus, the virtual circuit table 45 can effectively map a plurality of received virtual circuit identifier values to a single transmitted virtual circuit identifier value.

Thus, for example, with reference to network 10 (FIG. 1), since, in the illustrative virtual circuit described above in connection with that FIG., the virtual circuit branches at switching node 11(3), the virtual circuit table 45 for switching node 11(3) can have separate entries for each of the branches. In that case, the switching node 11(3)'s virtual circuit table 45 will include one entry which associates the received virtual circuit identifier value for virtual circuit packets that it receives over communication link 13(13) from switching node 11(4) to the transmitted virtual circuit identifier value that it uses in connection with transfer of virtual circuit packets over communication link 13(9) to the switching node 11(1). In addition, the switching node 11(3)'s virtual circuit table 45 will include a second entry which associates the received virtual circuit identifier value for virtual circuit packets that it receives over communication link 13(15) from switching node 11(5) to the same transmitted virtual circuit identifier value.

It will be appreciated that, by allowing an switching node 11(n) to make use of diverse virtual circuit identifier values for each respective virtual circuit in the manner described above, the virtual circuit identifier values for a virtual circuit need only be determined locally between neighboring switching nodes 11(n) and 11(n') (n'≠n) along the path for the virtual circuit, and not globally across the network 10. It will be appreciated that the actual values of the received and transmitted virtual circuit identifier values associated with a particular virtual circuit need not be unique, and, indeed, two or more of the three virtual circuit identifier values maintained by a switching node 11(n) for a particular virtual circuit may be the same.

The network topology database 46 contains information describing the topology of the network 10, which the control subsystem 41 can use in establishing the tables 44 and 45 and in controlling the establishment and extension of the virtual circuit branches for its particular virtual circuit (that is, the virtual circuit for which the switching node 11(n) is the terminating switching node 11(n)) over the network 10. The control subsystem 41 can update the information in the

network topology database 46 in response to advertising message packets which each of the switching nodes 11(n) generates and broadcasts throughout the network 10. Each advertising message packet provides local network topology information associated with the switching node that generated the advertising message packet, in particular identifying each the communication link 13(p) that is connected to the switching node that generated the advertising message packet and the other switching node that is connected thereto.

Using the network topology as defined by the network topology database 46, the control subsystem 41 can determine optimal paths for packet transfers to and from each of the other switching nodes 11(n') ($n' \neq n$) using operations defined by conventional optimal- or shortest-path algorithms. The control subsystem 41, for each switching node 11(n'), can store the information defining the path that it determines in the routing table 44. In addition, the control subsystem 41 can use the paths that it determines to identify, for each switching node 11(n') ($n' \neq n$), the path for its (the switching node 11(n)'s) virtual circuit from the respective switching node 11(n') to the particular switching node 11(n).

After the control subsystem 41 of switching node 11(n) determines the paths to each of the respective switching nodes 11(n'), it (that is, the control subsystem 41) will proceed to establish the path or paths for its virtual circuit. In that operation, the control subsystem 41 will initiate operations similar to those conventionally performed to establish "point to multi-point" virtual circuits. As is conventional, a point to multi-point virtual circuit is in the shape of a tree having at least one branch extending from a "root" switching node to a plurality of switching nodes, at least one of which is a "leaf" switching nodes. The "root" switching node, as originating switching node on a point to multi-point virtual circuit, uses the point to multi-point virtual circuit to multi-cast virtual circuit message packets to a plurality of terminating switching nodes, including the leaf switching nodes. To establish a virtual circuit in connection with the invention, the switching node 11(n), as terminating switching node, and other switching nodes 11(n') ($n' \neq n$) as originating switching nodes 11(n₀) or intermediate switching nodes 11(n_i), will perform operations corresponding to those performed by switching nodes in establishing a point to multi-point virtual circuit, except that:

(i) Each switching node 11(n') that can operate as an originating switching node 11(n₀) along the virtual circuit will generate an entry for its routing table 45 that contains, associated with the network address for the switching node 11(n) as terminating switching node 11(T), the virtual circuit identifier received from the previous switching node 11(n") ($n'' \neq n$) along the path of the virtual circuit (that is, the switching node 11(n") toward the terminating switching node 11(n_T) along the

-20-

virtual circuit) as the transmitted virtual circuit identifier value and the communication link 13(p) interconnecting the switching nodes 11(n') and 11(n'') as the communication link over which virtual circuit packets for the virtual circuit are to be transmitted. As indicated above, the switching node 11(n'), when it receives a device packet (FIG. 2) for transfer to the switching node 11(n), will use the information from the entry of its routing table in generating a virtual circuit packet (FIG. 3) and transmitting it to the next switching node 11(n'').

(ii) Each switching node 11(n') will in addition establish one or more entries in its virtual circuit table 45, such that

- (a) the communication link 13(p) interconnecting it (that is, switching node 11(n')) and the previous switching node 11(n'') ($n'' \neq n'$, n) along the virtual circuit path will be identified as the communication link over which virtual circuit packets associated with the virtual circuit are to be transmitted,
- (ii) the virtual circuit identifier value that is received from the previous switching node 11(n'') ($n'' \neq n'$, n) will be identified as the transmitted virtual circuit identifier value in the entry of the virtual circuit table 45 associated with the virtual circuit, and
- (iii) the virtual circuit identifier value that it (that is switching node 11(n')) uses in connection with next switching node 11(n'') ($n'' \neq n'$, n' , n) (if any) along the path to the leaf switching node will be identified as the received virtual circuit identifier value,

thereby to facilitate association of each virtual circuit packet received from the next switching node 11(n'') to a particular virtual circuit and communication link 13(p) over which it (that is, switching node 11(n')) is to transmit the virtual circuit packet. It will be appreciated that, if the switching node 11(n') is the remote switching node 11(n_R) along a particular path or branch of the virtual circuit, no received virtual circuit identifier value need be provided for the entry associated with the virtual circuit in its virtual circuit table 45. The effect of the above-described operations is to effectively establish the virtual circuit as a "multi-point to point" virtual circuit, with the root switching node as the terminating switching node 11(n_T) for the virtual circuit and the other switching nodes as originating switching nodes 11(n_O).

After the virtual circuit paths have been initially established, the switching node 11(n) will generally update the virtual circuit paths when it receives an advertising message packet that indicates that, for example, a new switching node has been added to the network 10, in which case a virtual circuit path will be established thereto, or a communication link 13(p) has malfunctioned,

in which case virtual circuit paths to the various switching nodes beyond the malfunctioning communication link 13(p) will need to be re-established to over another communication link. The switching nodes 11(n) can perform operations corresponding to those performed by switching nodes in updating a point to multi-point virtual circuit, modified as described above in connection with virtual circuit establishment.

With this background, selected operations performed by a switching node 11(n) in connection with the invention will be described in connection with the flow chart in FIG. 5, in particular operations performed by the switching nodes 11(n) in connection with establishing a virtual circuit. Preliminarily, operations performed by the switching nodes 11(n) in connection with

(i) receipt of device packets 20 and generation of respective virtual circuit packets 30 therefrom,

(ii) transfer of virtual circuit packets, both those generated by the switching node 11(n) and those received from other switching nodes, over the virtual circuit, and

(iii) extraction of devices packets 20 from virtual circuit packets 30 for delivery to a computer system 12(m),

are similar to those described in the aforementioned Chen application (except that in item (i) above the switching nodes 11(n) described herein use the routing table 44 to identify the particular virtual circuit to be used to transfer the virtual circuit packet 30 to the terminating switching node 11(n_T)), and will not be described herein.

FIG. 5 specifically depicts the operations that are performed by each switching node in connection with receipt of a virtual circuit path establishment message packet, which each terminating switching node 11(n_T) initially generates and transfers to enable establishment of respective paths for its virtual circuit. A virtual circuit path establishment message packet includes a virtual circuit identifier and a series of communication link identifiers. In one embodiment, in which virtual circuit identifier values are determined locally between neighboring switching nodes, the virtual circuit identifier in the virtual circuit path establishment message packet functions as a local virtual circuit identifier value. The series of communication link identifiers in the virtual circuit path establishment message packet define the series of communication links 13(p₁), 13(p₂),..., over which successive switching nodes are to transfer the virtual circuit path establishment message packet, thereby to establish the virtual circuit path. Operations performed by the terminating switching node 11(n_T) herein in connection with generation of a virtual circuit path establishment message packet are similar to those performed by an originating switching node 11(n₀) in connection

-22-

with generating a corresponding packet for generating a point to multi-point virtual circuit, and will not be described herein in detail.

With reference to FIG. 5, when a switching node 11(n) receives a virtual circuit path establishment message packet over a communication link 13(p_R) from another switching node 11(n') (n'≠n) (step 100), it (that is, the switching node 11(n)) will initially determine whether the entry in its routing table 44 associated with the terminating switching node 11(n_T) for the virtual circuit, which switching node had initially generated the virtual circuit path establishment message packet, includes a virtual circuit identifier that it (that is, the switching node 11(n)) can use as an originating switching node. If the switching node 11(n) makes a negative determination in step 101, it loads the virtual circuit identifier value from the virtual circuit path establishment message packet in the routing table entry. Following step 102, or if it makes a positive determination in step 101, or following step 102, the switching node 11(n) it will determine whether the virtual circuit path establishment message packet identifies a communication link 13(p_X) over which it is to transfer the virtual circuit path establishment message packet (step 103), and if not will generate an acknowledgment message for transmission to the terminating switching node 11(n_T) which originally generated the virtual circuit path establishment message packet to indicate that the virtual circuit path has been established (step 104).

Returning to step 103, if the switching node 11(n) makes a positive determination in that step (which will occur if the switching node 11(n) determines that the virtual circuit path establishment message packet identifies a communication link 13(p_X) over which it is to transfer the virtual circuit path establishment message), it will determine whether the virtual circuit table 45 includes an entry for the identified communication link 13(p_X) that is associated with the virtual circuit (step 110). If the switching node 11(n) makes a negative determination in step 110, it will establish such an entry (step 111), and load in the entry:

(i) an identifier for the communication link 13(p_R) over which it (that is, the switching node 11(n)) received the virtual circuit path establishment message packet as the communication link over which virtual circuit packets received by the switching node 11(n) for the virtual circuit are to be transmitted (step 112), and

(ii) the virtual circuit identifier from the received virtual circuit path establishment message packet as the transmitted virtual circuit identifier value (step 113).

In addition, the switching node 11(n) selects a virtual circuit identifier value to provide a local identification for the virtual circuit with the next switching node 11(n") (n"≠n) attached to the

-23-

communication link 13(p_X) over which it (that is, the switching node 11(n)) is to transmit the virtual circuit path establishment message packet (step 114), and load the selected virtual circuit identifier value in the entry of the virtual circuit table 45 as the received virtual circuit identifier value (step 115). The switching node 11(n) also loads the selected virtual circuit identifier value in the virtual circuit path establishment message packet as the local virtual circuit identifier value (step 116) and transmits the virtual circuit path establishment message packet to the next switching node 11(n") over the communication link 13(p_X) (step 117).

Returning to step 110, if the switching node 11(n) makes a positive determination in that step, which will occur if the virtual circuit table 45 already contains an entry for the virtual circuit for the communication link 13(p_X), it (that is, the switching node 11(n)) loads the received virtual circuit identifier value from the virtual circuit table entry into the virtual circuit path establishment message packet (step 120) as the local virtual circuit identifier value, and transfers the virtual circuit path establishment message packet to the next switching node 11(n") over the communication link 13(p_X) (step 121).

It will be appreciated that a number of modifications may be made to the network 10 as described above in connection with FIGS. 1 through 5. For example, although the invention has been described such that the terminating switching node 11(n_T) controls establishment of the virtual circuit paths defining its virtual circuit, it will be appreciated that other switching nodes, as remote switching nodes 11(n_R), along each path can instead control establishment of the virtual circuit paths to each respective terminating switching node 11(n_T).

Furthermore, although the invention has been described in connection with use of data structures, including the a routing table 44, virtual circuit table 45, and network topology database 46 described herein and in the aforementioned Chen application, it will be appreciated that data structures may be used having a variety of structures and organizations instead of and/or in addition for providing the types of information as described above.

It will be appreciated that a system in accordance with the invention can be constructed in whole or in part from special purpose hardware or a general purpose computer system, or any combination thereof, any portion of which may be controlled by a suitable program.

The foregoing description has been limited to a specific embodiment of this invention. It will be apparent, however, that various variations and modifications may be made to the invention, with the attainment of some or all of the advantages of the invention. It is the object of the appended

-24-

claims to cover these and such other variations and modifications as come within the true spirit and scope of the invention.

What is claimed as new and desired to be secured by Letters Patent of the United States is:

CLAIMS

1. A network comprising a plurality of switching nodes interconnected by a plurality of communication links, the switching nodes transferring message packets thereamong, at least one virtual circuit being established in said network to at least one said switching node extending over a plurality of others of said switching nodes, the virtual circuit defining a path over which said others of said switching nodes transfer ones of said message packets to said at least one switching node.
2. A network as defined in claim 1 in which said path extends over all of said others of said switching nodes, all of said others of said switching nodes using said virtual circuit to transfer ones of said message packets to said at least one said switching node.
3. A network as defined in claim 1 in which said virtual circuit comprises a plurality of paths extending to said at least one said switching node extending over a plurality of others of said switching nodes.
4. A network as defined in claim 1 in which said at least one virtual circuit path is established in response to a virtual circuit path establishment message packet generated by said at least one switching node.
5. A network as defined in claim 4 in which:
 - A. said virtual circuit path establishment message packet includes a virtual circuit identifier value identifying said virtual circuit; and
 - B. said at least one of said others of said switching nodes includes a router table including a router table entry which associates said at least one switching node and a virtual circuit identifier, said at least one of said others of said switching nodes using the virtual circuit identifier value as the virtual circuit identifier for said router table entry.
6. A network as defined in claim 5 in which said at least one of said others of said switching nodes includes a virtual circuit packet generator for generating virtual circuit packets for transferring information to said at least one switching node, each virtual circuit packet including a virtual circuit packet identifier field for receiving a virtual circuit identifier that identifies the virtual circuit to be used in transferring the virtual circuit packet, said virtual circuit packet generator using the virtual circuit identifier value from said router table entry as the virtual circuit identifier when generating a virtual circuit packet for use in transferring information to said one of said switching nodes.
7. A network as defined in claim 6 in which:
 - A. said virtual circuit path establishment message packet further includes a virtual circuit identifier value for identifying said virtual circuit;

-26-

- B. said at least one of said others of said switching nodes includes:
 - (i) a virtual circuit table; and
 - (ii) a control subsystem responsive to said virtual circuit path establishment message packet for selectively establishing an entry in said virtual circuit table, the entry identifying said virtual circuit identifier value and the one of said communication links over which said switching node received said virtual circuit path establishment message packet.
- 8. A network as defined in claim 7 in which said control subsystem uses the virtual circuit identifier value in connection with transfers of message packets that said at least one of said others of said switching nodes receives for the virtual circuit.
- 9. A network as defined in claim 7 in which:
 - A. said at least one of said others of said switching nodes is connected to a plurality of said communication links, each of said plurality of said communication links being identified by a communication link identifier;
 - B. said virtual circuit path establishment message packet includes a communication link identifier that identifies one of said plurality of said communication links over which said at least one of said others of said switching nodes is to transfer the virtual circuit path establishment message packet; and
 - C. said at least one of said others of said switching nodes includes a packet switch for transferring said virtual circuit path establishment message packet over said one of said plurality of said communication links identified by said communication link identifier.
- 10. A network as defined in claim 9 in which said control subsystem generates a received virtual circuit identifier value and inserts the received virtual circuit identifier value in the virtual circuit path establishment message packet prior the packet switch transfer, the control subsystem further using the received virtual circuit identifier value in the virtual circuit entry.
- 11. A network as defined in claim 10 in which said control subsystem uses the received virtual circuit identifier value in connection with each said message packet received thereby to identify the virtual circuit associated therewith.
- 12. A method of controlling a network comprising a plurality of switching nodes interconnected by communication links, the switching nodes transferring message packets thereamong, the method comprising the steps of

-27-

- A. establishing at least one virtual circuit in said network to at least one said switching node extending over a plurality of others of said switching nodes, said at least one virtual circuit comprising a , and
- B. enabling others of said switching nodes to use the virtual circuit defining a path over which said others of said switching nodes transfer ones of said message packets to said at least one switching node.

13. A network comprising a plurality of switching nodes interconnected by a plurality of communication links, the switching nodes transferring message packets thereamong, at least one virtual circuit being established in said network defining a plurality of paths to a like plurality of others of said switching nodes which said others of said switching nodes transfer ones of said message packets to said at least one switching node.

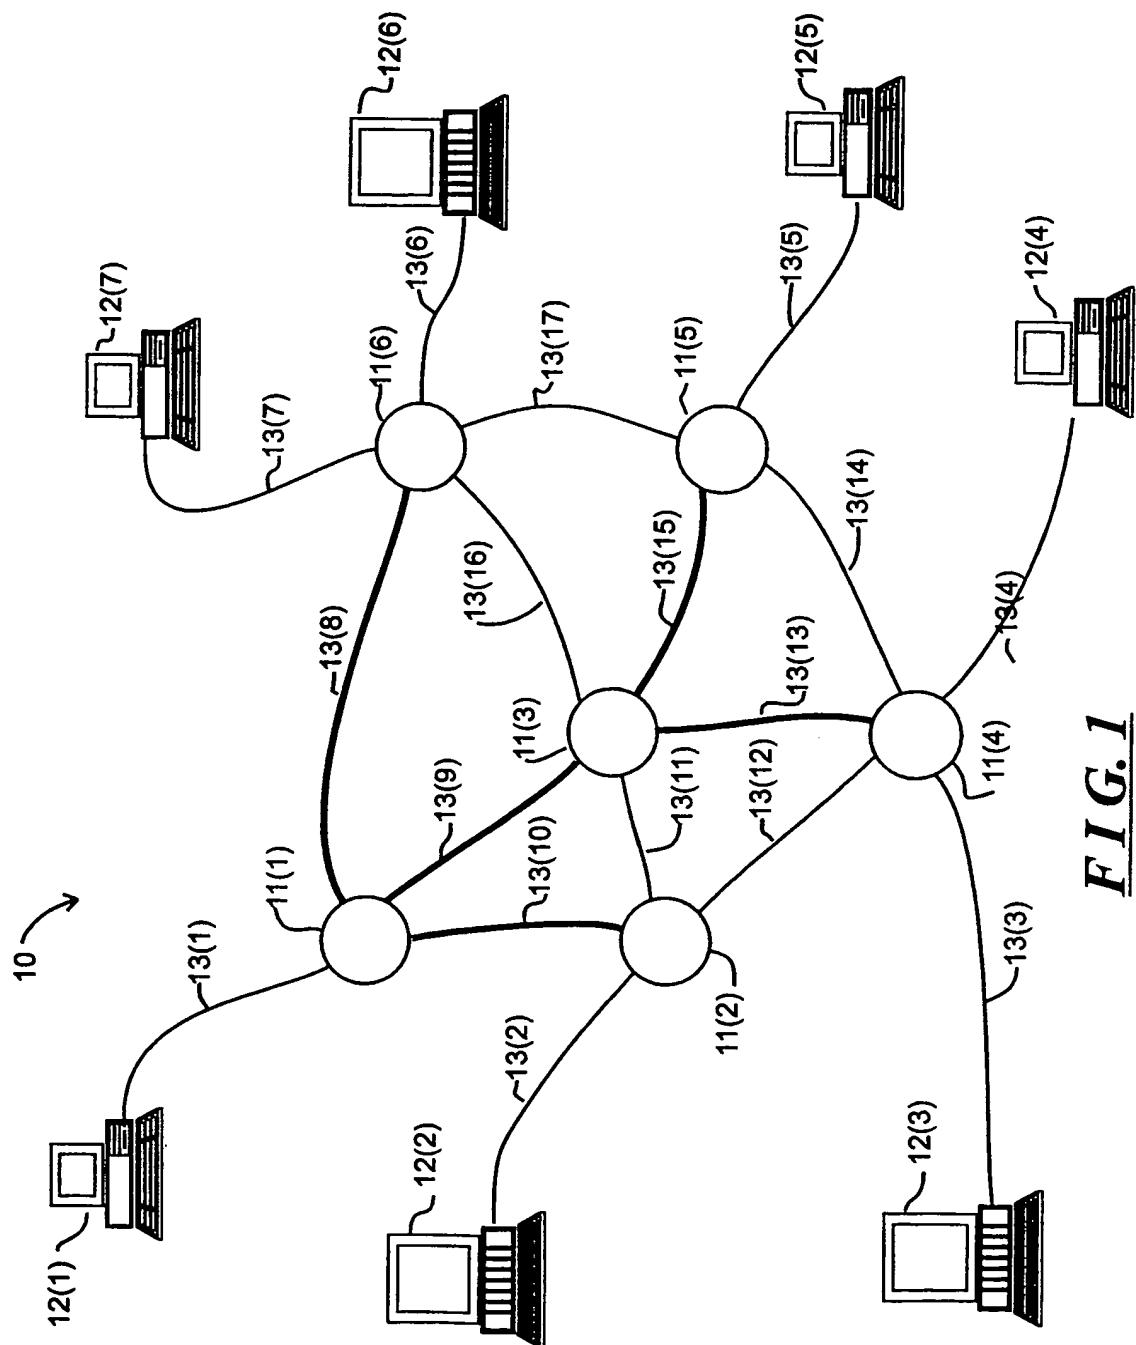
14. A switching node for use in a network comprising:

- A. a switching fabric for receiving and transmitting message packets over a plurality of communication links, said message packets including:
 - i. message packets of a virtual circuit path establishment message packet type, each message packet of the virtual circuit path establishment message packet type including a virtual circuit identifier value and a next communication link identifier;
 - ii. message packets of a virtual circuit packet type, each virtual circuit packet including a received virtual circuit identifier;
- B. a virtual circuit table; and
- C. a control subsystem for:
 - i. in response to a message packet of the virtual circuit path establishment message packet type,
 - a. selectively establishing an entry in said virtual circuit table, said entry including a communication link identifier for identifying one of said plurality of communication links over which said virtual circuit path establishment message packet type and the virtual circuit identifier value as a transmitted virtual circuit identifier value,
 - b. selecting a received virtual circuit identifier value for said entry,
 - c. providing said received virtual circuit identifier value in said message packet of the virtual circuit path establishment message and transmitting the message packet of the virtual circuit path establishment message type over the one of

-28-

said plurality of communication links identified by said next communication link identifier; and

- ii. in response to a message packet of the virtual circuit packet type,
 - a. using the received virtual circuit identifier in the message packet of the virtual circuit packet type to identify a virtual circuit table entry,
 - b. providing the transmitted virtual circuit identifier value in the entry virtual circuit packet, and transmitting the message packet of the virtual circuit packet type over the one of said plurality of communication links identified by said entry.



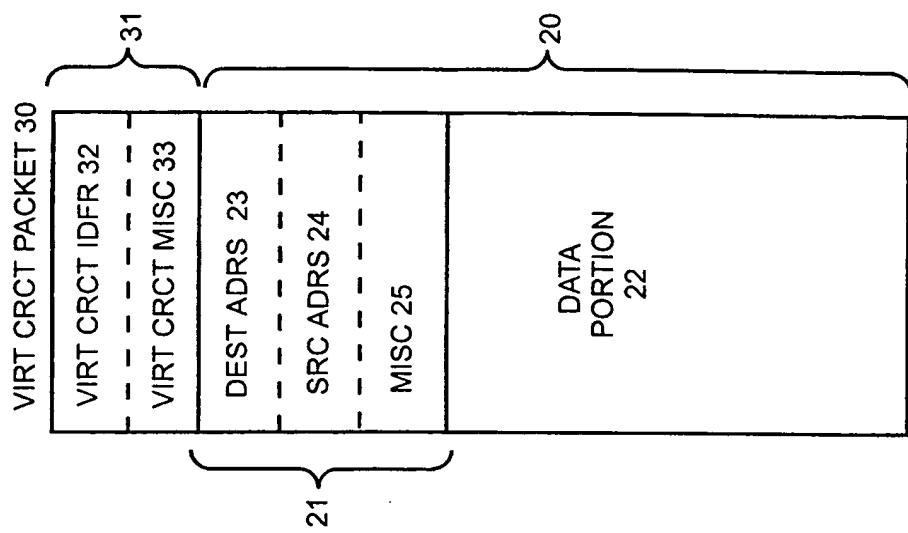


FIG. 3

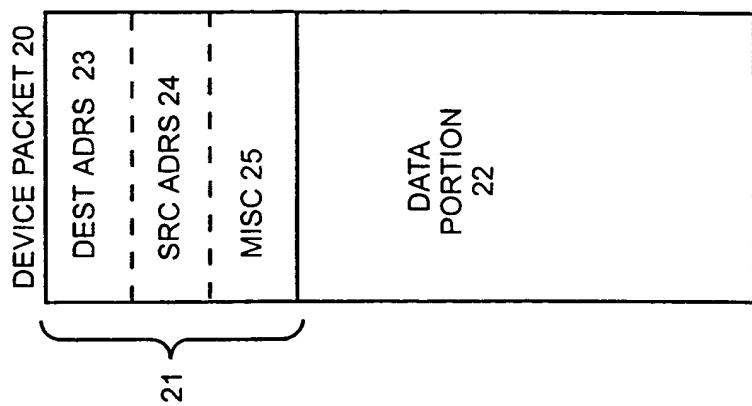


FIG. 2

3/9

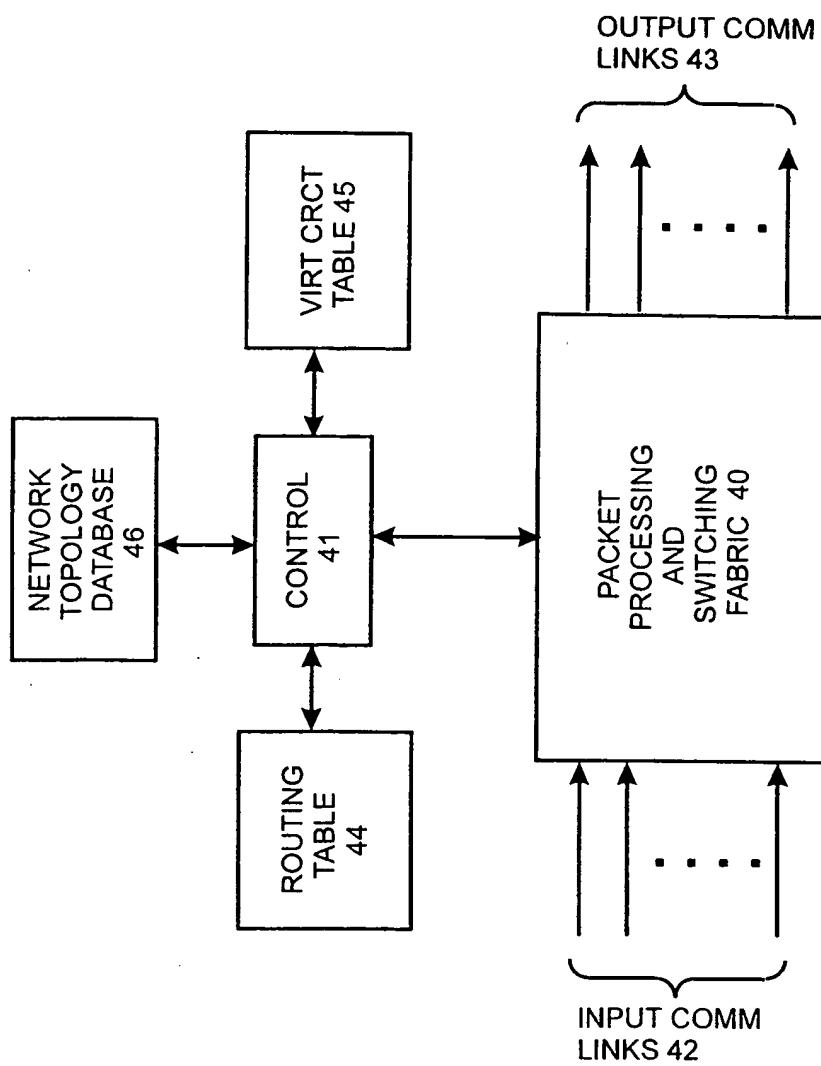
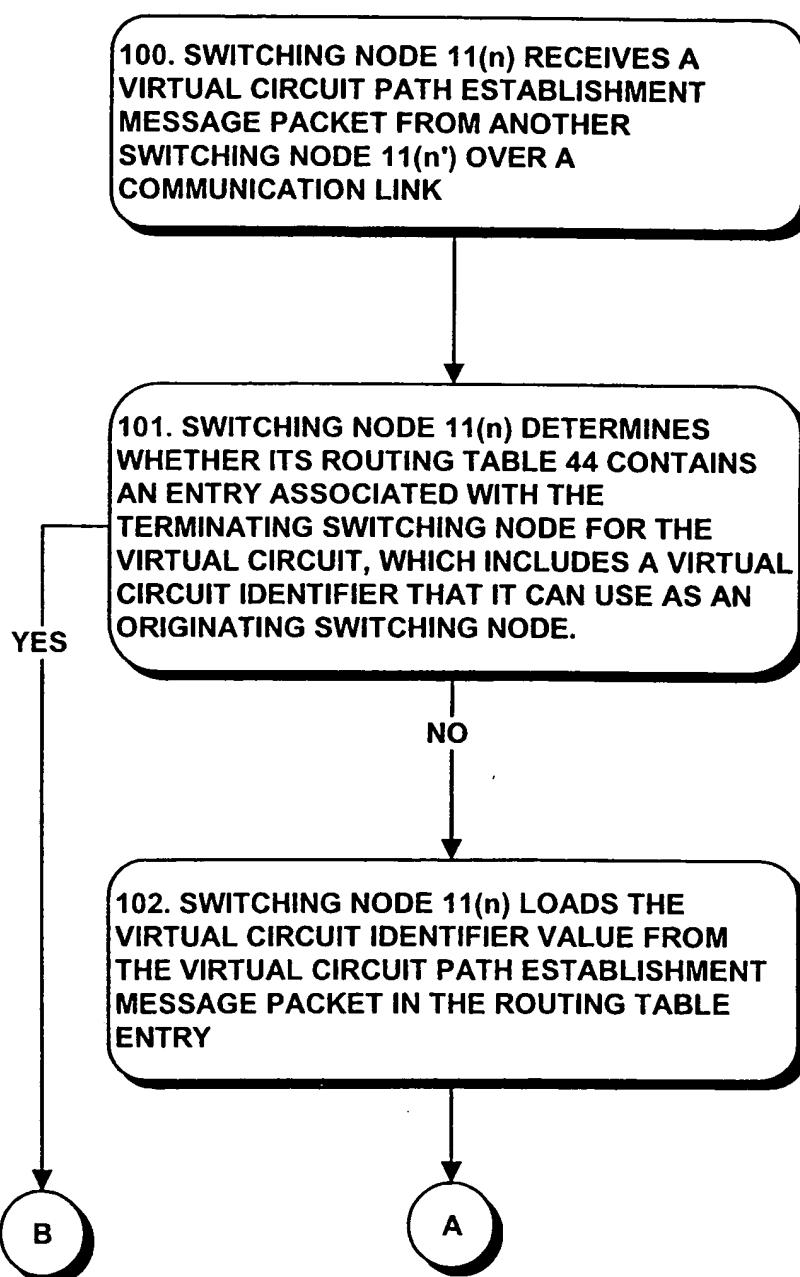
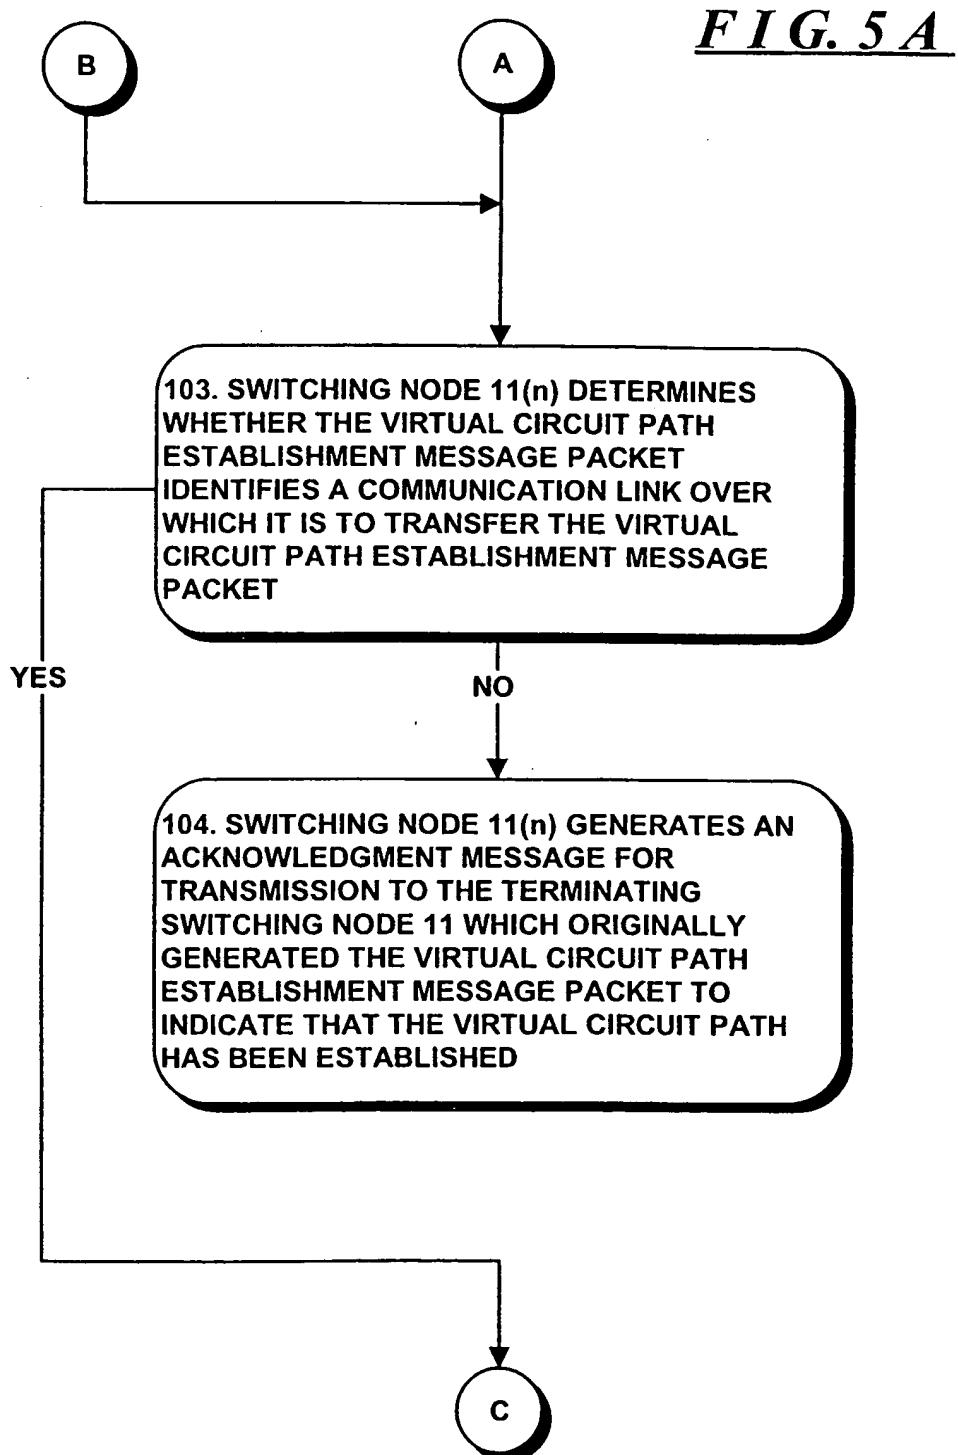


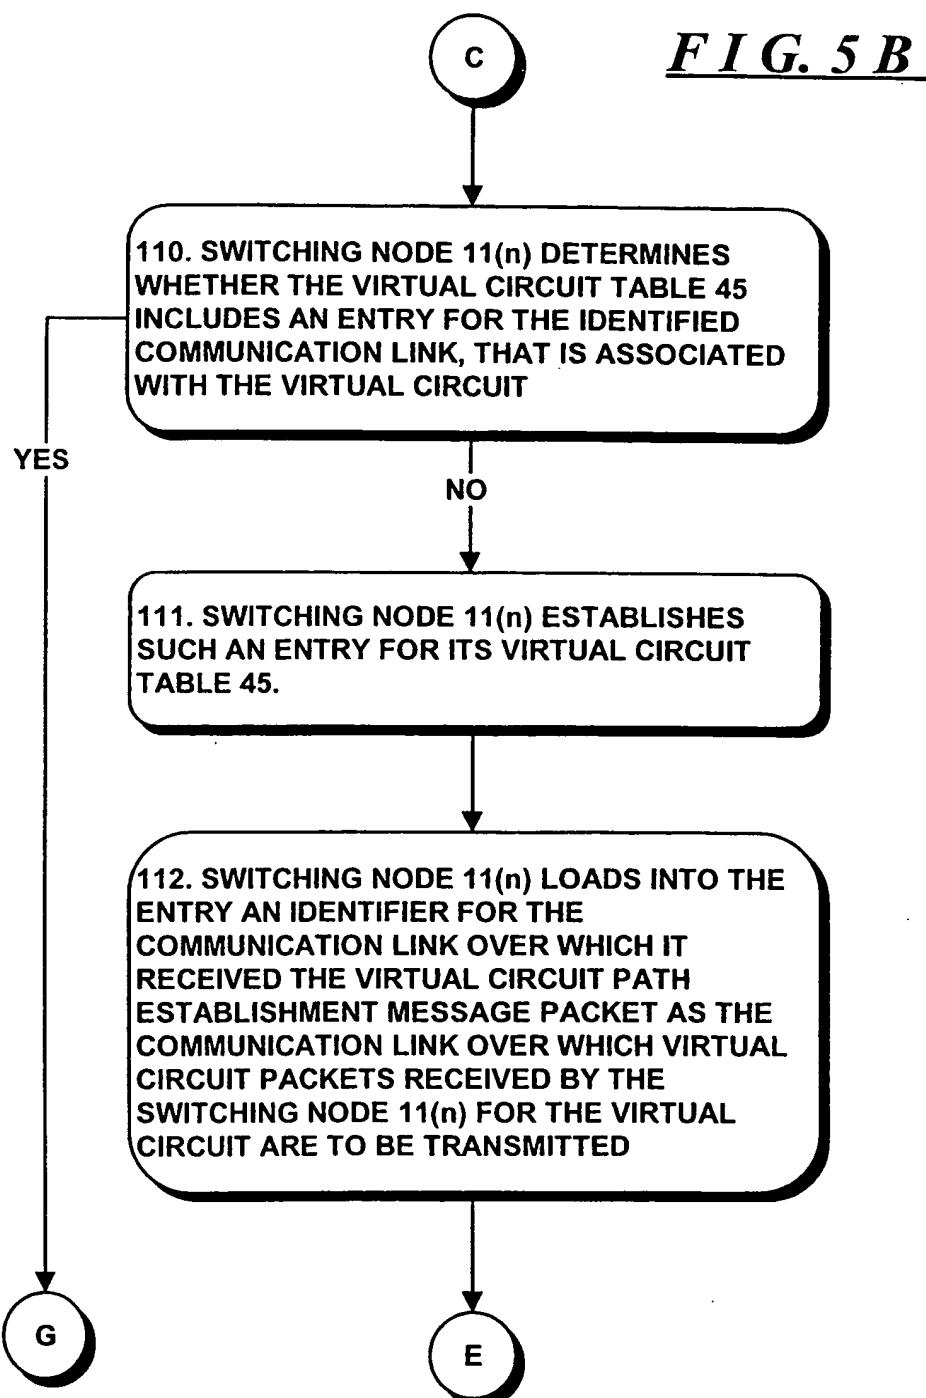
FIG. 4

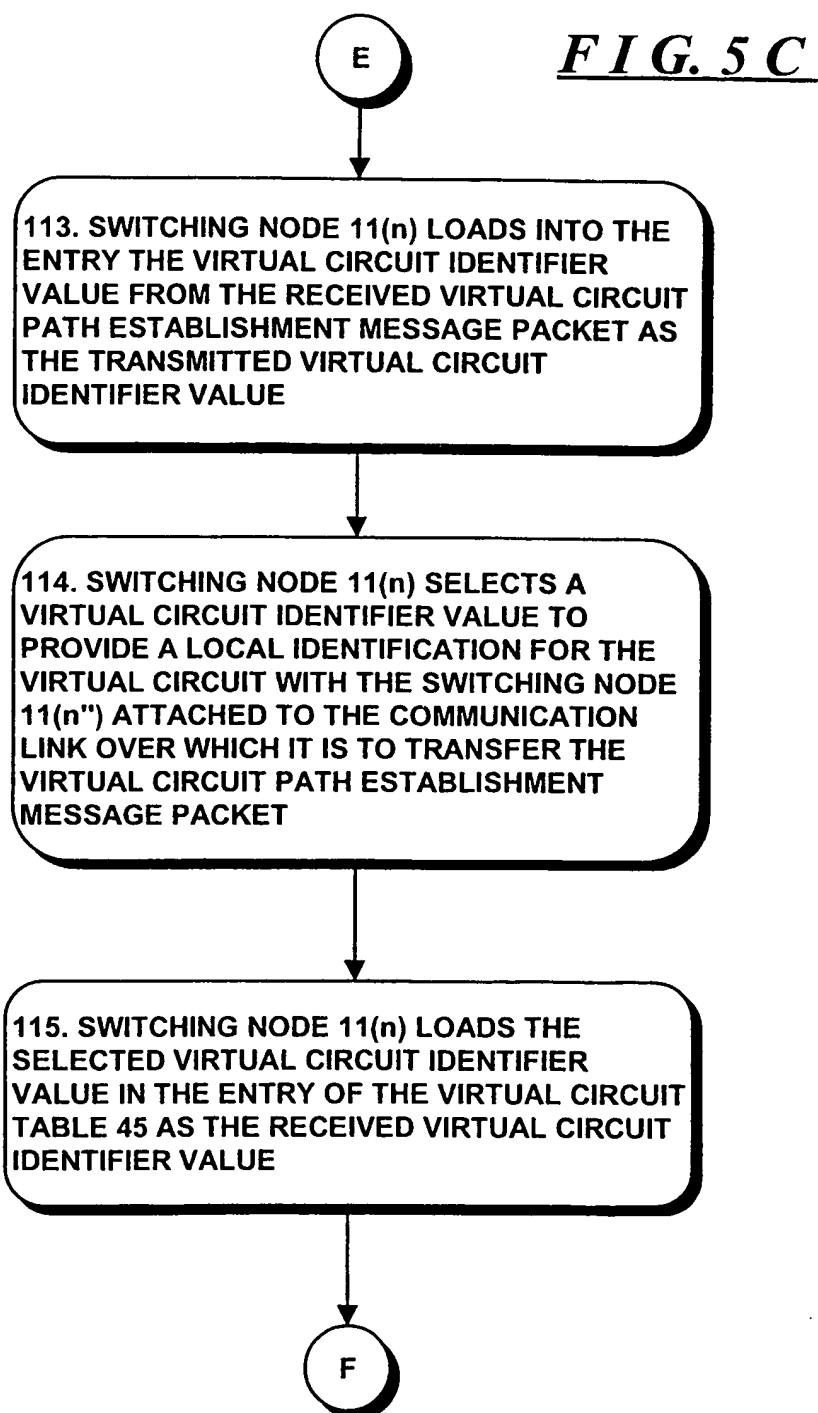
SWITCHING NODE 11(n)

FIG. 5

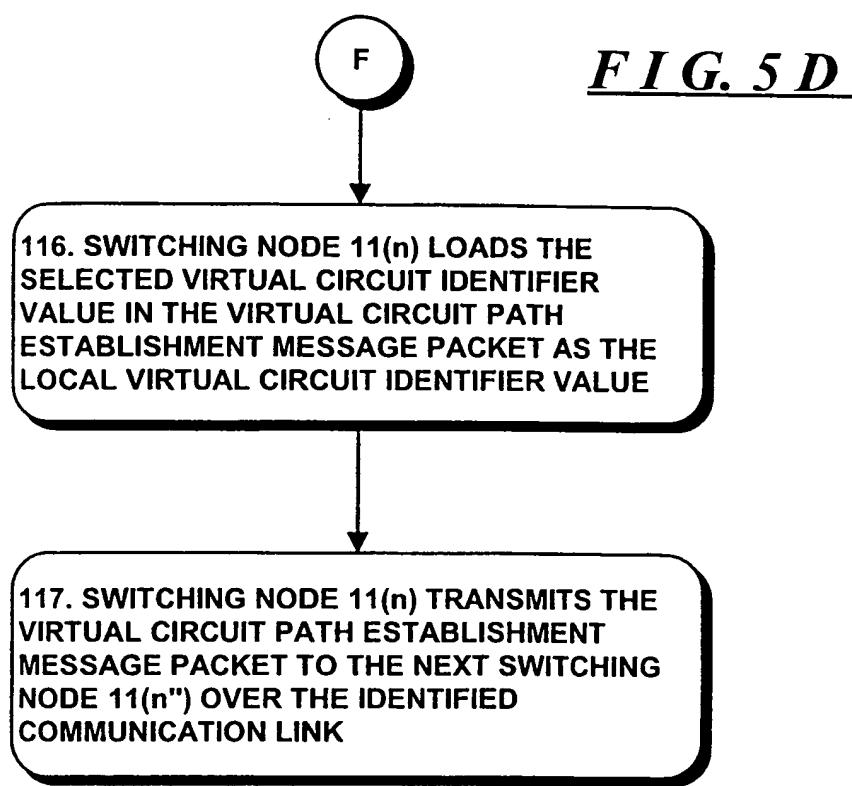


6/9

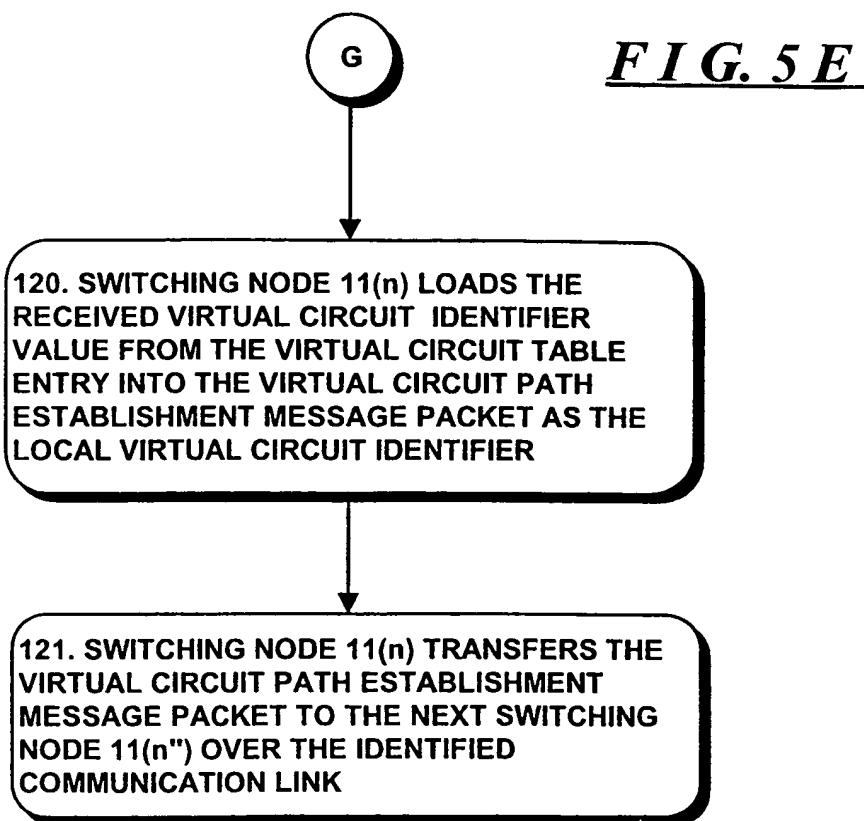




8/9



9/9



INTERNATIONAL SEARCH REPORT

International application No.

PCT/US97/19558

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) :H04L 12/46, 12/56

US CL :370/355, 397, 409

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 370/254, 255, 355, 356, 389, 395-401, 409

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

None

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

APS: (packet# or cell#), switch?(w)node#, virtual(w)(circuit or connection or channel), multipoint(2w)point

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5,357,508 A (LE BOUDEC et al) 18 October 1994, col. 6, line 33 to col. 7, line 26).	1-14
X	US 5,450,406 A (ESAKI et al) 12 September 1995, col. 16, line 9 to col. 17, line 68).	1-14
X,P	US 5,651,005 A (KWOK et al) 22 July 1997, abstract, col. 3, line 57 to col. 5, line 57).	1-14
A	US 5,440,547 A (ESAKI et al) 08 August 1995, see entire document.	1-14
A,P	US 5,633,866 A (CALLON) 27 May 1997, see entire document.	1-14

 Further documents are listed in the continuation of Box C. See patent family annex.

• Special categories of cited documents:	"T"	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
A document defining the general state of the art which is not considered to be of particular relevance	"X"	document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
B earlier document published on or after the international filing date	"Y"	document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"A"	document member of the same patent family
O document referring to an oral disclosure, use, exhibition or other means		
P document published prior to the international filing date but later than the priority date claimed		

Date of the actual completion of the international search

02 JANUARY 1998

Date of mailing of the international search report

25 FEB 1998

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

B. Hsu

ALPUS H. HSU

Telephone No. (703) 305-4377